

Preparing to Run on Aitken Cascade Lake Nodes

To help you prepare for running jobs on Aitken's Cascade Lake compute nodes, this short user guide includes information on the general configuration of Cascade Lake nodes, compiling your code, running PBS jobs, and checking allocation usage.

Overview of Aitken Cascade Lake Nodes

Aitken includes 1,152 Cascade Lake nodes, which are partitioned into eight physical racks. Each node contains two 20-core Xeon Gold 6248 sockets (2.5 GHz) and 192 GB of memory. The nodes are connected to the Aitken InfiniBand network (ib0 and ib1) via four-lane Enhanced Data Rate (4X EDR) devices and four-lane High Data Rate (4X HDR) switches for internode communication. The ib1 fabric is used mainly for I/O and is connected to the Pleiades Lustre filesystems. In addition, Aitken, Electra, and Pleiades share the same home filesystems, Pleiades front-end systems (PFEs), and PBS server. You can access Aitken only through PBS jobs submitted from the PFEs.

Compiling Your Code For Cascade Lake Nodes

The Cascade Lake processors include the Advanced Vector Extensions 512 (AVX-512). Intel AVX-512 optimizations are included in Intel compiler version 16.0 and later versions. We recommend that you test the latest Intel compiler, through the command `module load comp-intel/2020.4.304`, which may include better optimizations for AVX-512.

For Cascade Lake-specific optimizations, use the compiler option `-xCORE-AVX512`.

If you want a single executable that will run on any of the Aitken, Electra, and Pleiades processor types, with suitable optimization to be determined at runtime, you can compile your application using the option:

`-O3 -xCORE-AVX512, CORE-AVX2 -xAVX.`

Note: The use of either of these options, `-xCORE-AVX512` or `-xCORE-AVX512`, could either improve or degrade performance of your code. Be sure to check performance with and without these flags before using them for production runs.

Running PBS Jobs on Aitken Cascade Lake Nodes

To request Aitken Cascade Lake nodes, use `:model=cas_ait` in your PBS script, as shown in the example below.

Note: Because MPT 2.15 and earlier versions do not support the ConnectX-5 host channel adapters (HCAs), the environment variables `MPI_IB_XRC` and `MPI_XPMEM_ENABLED` have been disabled for jobs running on Cascade Lake. If your MPI applications perform significant MPI collective operations and rely on having these two variables enabled to get good performance, use MPT 2.17 or newer versions. You can use the NAS-recommended MPT version with the command:

```
module load mpi-hpe/mpt
```

Sample PBS Script For Aitken Cascade Lake Nodes

```
#PBS -l select=10:ncpus=40:mpiprocs=40:model=cas_ait
#PBS -l walltime=8:00:00
```

```
#PBS -q normal  
  
module load mpi-hpe/mpt  
module load comp-intel/2020.4.304  
  
cd $PBS_0_WORKDIR  
  
mpiexec -np 400 ./a.out
```

For more information, see [Cascade Lake Processors](#).

Article ID: 597
Last updated: 13 May, 2021
Revision: 11
Systems Reference -> Aitken -> Preparing to Run on Aitken Cascade Lake Nodes
<https://www.nas.nasa.gov/hecc/support/kb/entry/597/>